

WEST Search History

Hide Items

Restore

Clear

Cancel

DATE: Tuesday, February 28, 2006

Hide?	<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L59	L58 and checkpoint\$1	1
<input type="checkbox"/>	L58	(backup near\$ database\$1) and (rebuil\$ near\$ database\$1) and (transaction near\$ log\$1) and @py<=2003	3
<input type="checkbox"/>	L57	L56 and (fault near\$ tolerance)	2
<input type="checkbox"/>	L56	L55 and (transaction\$1 near\$ log\$1)	15
<input type="checkbox"/>	L55	(recover\$3 near\$ log\$1) and (backup near\$ database\$1) and checkpoint\$1 and @py<=2003	22
<input type="checkbox"/>	L54	(transaction\$1 near\$ log\$1) and (new near\$ checkpoint\$1) and (backup near\$ database\$1) and @py<=2003	1
<input type="checkbox"/>	L53	disk and backup and storage and checkpoint\$1 and transaction and log\$1 and updat\$3 and (id\$ or unique id\$) and database\$1 and (read near\$ log\$1) and (sort\$3 near\$ log\$1) and @py<=2003	0
<input type="checkbox"/>	L52	(fault near\$ tolerance) and backup and (transaction near\$ log\$1) and checkpoint\$1 and merg\$3 and @py<=2003	3
<input type="checkbox"/>	L51	(checkpoint\$1 near\$ interval\$1) and (transaction near\$ log\$1) and (rebuilt near\$ database\$1) and backup and merg\$3 and sort\$3 and updat\$3 and @py<=2003	0
<input type="checkbox"/>	L50	L49 and merg\$3	5
<input type="checkbox"/>	L49	L48 and checkpoint\$1	9
<input type="checkbox"/>	L48	L47 and (log near\$ id\$)	9
<input type="checkbox"/>	L47	L46 and (log near\$ record\$1)	28
<input type="checkbox"/>	L46	L45 and (transaction near\$ log\$1)	28
<input type="checkbox"/>	L45	L44 and (log near\$ record\$1)	28
<input type="checkbox"/>	L44	(database\$1 near\$ backup) and (transaction near\$ log\$1) and checkpoint\$1 and updat\$3 and @py<=2003	34
<input type="checkbox"/>	L43	(transact\$4 and log\$1 and database\$1 and backup\$).ti.	4
<input type="checkbox"/>	L42	L41 and updat\$3	4
<input type="checkbox"/>	L41	L40 and combin\$3	4
<input type="checkbox"/>	L40	L39 and stor\$3	4
<input type="checkbox"/>	L39	L38 and sort\$3	4
<input type="checkbox"/>	L38	L37 and merg\$3	5
<input type="checkbox"/>	L37	L36 and (log\$1 near\$ id\$)	8
<input type="checkbox"/>	L36	L35 and (checkpoint\$1 same database\$1)	20
<input type="checkbox"/>	L35	L34 and (transaction\$1 near\$ log\$1)	196

<input type="checkbox"/>	L34	(backup near5 database\$1) and @py<=2003	1686
<input type="checkbox"/>	L33	L32 and rebuilt	1
<input type="checkbox"/>	L32	L31 and checkpoint\$1	16
<input type="checkbox"/>	L31	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1) and (updat\$3 near5 database\$1) and @py<=2003	91
<input type="checkbox"/>	L30	L29 and (transaction near5 id\$) and checkpoint\$1 and (old near5 checkpoint\$1) and (new near5 checkpoint\$1) and merg\$3 and sort\$3 and stor\$3 and memory	0
<input type="checkbox"/>	L29	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1)	207
<input type="checkbox"/>	L28	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1) and (transaction near5 id\$) and checkpoint\$1 and (old near5 checkpoint\$1) and (new near5 checkpoint\$1) and merg\$3 and sort\$3 and stor\$3 and memory	0
<input type="checkbox"/>	L27	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1) and (transaction near5 id\$) and checkpoint\$1 and (old near5 checkpoint\$1) and (new near5 checkpoint\$1) and merg\$3 and sort\$3 and stor\$3 and memory and @py<=2003	0
<input type="checkbox"/>	L26	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1) and (transaction near5 id\$) and checkpoint\$1 and (old near5 checkpoint\$1) and (new near5 checkpoint\$1) and merg\$3 and sort\$3 and stor\$3 and memory and @py<=2003	0
<input type="checkbox"/>	L25	(backup near5 database\$1) and (log near5 transaction\$1) and (log near5 database\$1) and (transaction near5 id\$) and checkpoint\$1 and (old nar5 checkpoint\$1) and (new near5 checkpoint\$1) and merg\$3 and sort\$3 and stor\$3 and memory and @py<=2003	0
<input type="checkbox"/>	L24	L22 and (transaction near5 id\$)	4
<input type="checkbox"/>	L23	L22 and merg\$3 and sort\$3	0
<input type="checkbox"/>	L22	(database\$1 and backup and log\$1).ti,ab. and @py<=2003	64
<input type="checkbox"/>	L21	L20 and (transaction near5 log\$1)	4
<input type="checkbox"/>	L20	L18 and (log near5 id\$)	14
<input type="checkbox"/>	L19	L18 and (log near5 id\$)	14
<input type="checkbox"/>	L18	database\$1 and backup and transaction\$1 and checkpoint\$1 and merg\$3 and sort\$3 and stor\$3 and log\$1 and generat\$3 and @py<=2003	60
<input type="checkbox"/>	L17	database\$1 and backup and transaction\$1 and checkpoint\$1 and merg\$3 and sort\$3 and stor\$3 and @py<=2003	61
<input type="checkbox"/>	L16	l7 and log\$1 and id\$ and merg\$3 and sort\$3 and transaction\$1 and checkpoint\$1 and stor\$3 and memory	0
<input type="checkbox"/>	L15	l7 and (merg\$3 near5 log\$1)	1
<input type="checkbox"/>	L14	L12 and sort\$3	0
<input type="checkbox"/>	L13	L12 and merg\$3	0
<input type="checkbox"/>	L12	L11 and checkpoint\$1	1
<input type="checkbox"/>	L11	L10 and (log near5 id\$1)	1
<input type="checkbox"/>	L10	l7 and (transaction near5 id\$1)	12

<input type="checkbox"/>	L9	L8 and checkpoint\$1	3
<input type="checkbox"/>	L8	(backup and database\$1 and log\$1).ti,ab. and @py<=2002	46
<input type="checkbox"/>	L7	(backup and database\$1).ti,ab. and @py<=2002	486
<input type="checkbox"/>	L6	(backup and log\$1).ti. and @py<=2002	16
<input type="checkbox"/>	L5	(backup and transaction\$1 and log\$1).ti.	4
<input type="checkbox"/>	L4	L3 and merg\$3	2
<input type="checkbox"/>	L3	(database\$1 and backup and log\$1 and transact\$4).ti,ab. and @py<=2002	23
<input type="checkbox"/>	L2	(database\$1 and backup and log\$1 and transact\$4 and merg\$3).ti,ab. and @py<=2002	0
<input type="checkbox"/>	L1	(backup and log and record\$1 and transaction\$1 and checkpoint\$1 and (log near5 id\$) and database\$1 and data and handl\$3 and transmit\$3 and stack\$1 and memory and storage and generat\$3 and receiv\$3) and @py<=2002	6

END OF SEARCH HISTORY

WEST Search History

[Hide Items](#)
[Restore](#)
[Clear](#)
[Cancel](#)

DATE: Tuesday, February 28, 2006

Hide?	<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L16	(backup near5 checkpoint\$1) and (transaction near5 log\$1) and @py<=2003	16
<input type="checkbox"/>	L15	L14 and recover\$3	6
<input type="checkbox"/>	L14	(checkpoint near5 merg\$3) and @py<=2003	18
<input type="checkbox"/>	L13	L11 and (merg\$3 near5 checkpoint\$1)	0
<input type="checkbox"/>	L12	L11 and (updat\$3 near5 checkpoint\$1)	0
<input type="checkbox"/>	L11	L10 and (log near5 id\$)	14
<input type="checkbox"/>	L10	database\$1 and backup and transaction\$1 and log\$1 and checkpoint\$1 and merg\$3 and sort\$3 and stor\$3 and @py<=2003	60
<input type="checkbox"/>	L9	(checkpoint\$1 near5 merg\$3) and (checkpoint\$1 near5 sort\$3)	2
<input type="checkbox"/>	L8	L6 and (checkpoint\$1 near5 sort\$3)	0
<input type="checkbox"/>	L7	L6 and (checkpoint\$1 near5 merg\$3)	0
<input type="checkbox"/>	L6	L5 and (transaction near5 log\$1)	28
<input type="checkbox"/>	L5	(database\$1 same checkpoint\$1) and (backup same log\$1) and @py<=2003	50
<input type="checkbox"/>	L4	L2 and updat\$3	14
<input type="checkbox"/>	L3	L2 and checkpoint\$1	1
<input type="checkbox"/>	L2	(backup near5 transaction\$1) and (transaction near5 log\$1) and sort\$3 and merg\$3 and @py<=2003	14
<input type="checkbox"/>	L1	(generat\$3 near5 checkpoint\$1) and (merg\$3 near5 checkpoint\$1) and (stor\$3 near5 checkpoint\$1) and (transaction near5 log\$1) and @py<=2003	1

END OF SEARCH HISTORY


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[backup databases transaction log identifiers checkpoints](#)

Found 25,433 of 171,143

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [ARIES: a transaction recovery method supporting fine-granularity locking and partial](#)


[rollbacks using write-ahead logging](#)

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(5.23 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

2 [System R: relational approach to database management](#)



M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade, V. Watson

June 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available: pdf(3.18 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

System R is a database management system which provides a high level relational data interface. The systems provides a high level of data independence by isolating the end user as much as possible from underlying storage structures. The system permits definition of a variety of relational views on common underlying data. Data control features are provided, including authorization, integrity assertions, triggered transactions, a logging and recovery subsystem, and facilities for maintaining ...

Keywords: authorization, data structures, database, index structures, locking, nonprocedural language, recovery, relational model

3 Principles of transaction-oriented database recovery



Theo Haerder, Andreas Reuter

December 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available: pdf(2.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

4 High speed on-line backup when using logical log operations



David B. Lomet

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2

Publisher: ACM Press

Full text available: pdf(220.69 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Media recovery protects a database from failures of the stable medium by maintaining an extra copy of the database, called the backup, and a media recovery log. When a failure occurs, the database is "restored" from the backup, and the media recovery log is used to roll forward the database to the desired time, usually the current time. Backup must be both fast and "on-line", i.e. concurrent with on-going update activity. Conventional online backup sequentially copies ...

5 Recovery Techniques for Database Systems



Joost S. M. Verhofstad

June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Publisher: ACM Press

Full text available: pdf(2.32 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 The Recovery Manager of the System R Database Manager



Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available: pdf(1.75 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 Practical byzantine fault tolerance and proactive recovery



Miguel Castro, Barbara Liskov

November 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 4

Publisher: ACM Press

Full text available: pdf(1.63 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

8 Distributed, object-based programming systems



Roger S. Chin, Samuel T. Chanson

March 1991 **ACM Computing Surveys (CSUR)**, Volume 23 Issue 1

Publisher: ACM Press

Full text available: pdf(2.97 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The development of distributed operating systems and object-based programming languages makes possible an environment in which programs consisting of a set of interacting modules, or objects, may execute concurrently on a collection of loosely coupled processors. An object-based programming language encourages a methodology for designing and creating a program as a set of autonomous components, whereas a distributed operating system permits a collection of workstations or personal computers ...

Keywords: capability scheme, distributed operating systems, error recovery, method invocation, nested transaction, object model, object reliability, object-based programming languages, processor allocation, resource management, synchronization, transaction

9 A practical guide to the design of differential files for recovery of on-line databases



Houtan Aghili

December 1982 **ACM Transactions on Database Systems (TODS)**, Volume 7 Issue 4

Publisher: ACM Press

Full text available: pdf(1.54 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The concept of a differential file has previously been proposed as an efficient means of collecting database updates for on-line systems. This paper studies the problem of database backup and recovery for such systems, and presents an analytic model of their operation. Five key design decisions are identified and an optimization procedure for each is developed. A design algorithm that quickly provides parameters for a near-optimal differential file architecture is provided.

Keywords: backup and recovery, database maintenance, differential files, hashing functions, numerical methods, optimization, reorganization

10 On the selection of efficient record segmentations and backup strategies for large shared databases



Salvatore T. March, Gary D. Scudder

September 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 3

Publisher: ACM Press

Full text available: pdf(1.97 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the information processing requirements of business organizations have expanded tremendously. With this expansion, the design of databases to efficiently manage and protect business information has become critical. We analyze the impacts of record segmentation (the assignment of data items to segments defining subfiles), an efficiency-oriented design technique, and of backup and recovery strategies, a data protection technique, on the overall ...

11 A database cache for high performance and fast restart in database systems



Klaus Elhardt, Rudolf Bayer

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Publisher: ACM Press

Full text available: [pdf\(1.72 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Performance in database systems is strongly influenced by buffer management and transaction recovery methods. This paper presents the principles of the database cache, which replaces the traditional buffer. In comparison to buffer management, cache management is more carefully coordinated with transaction management, and integrates transaction recovery. High throughput of small- and medium-sized transactions is achieved by fast commit processing and low database traffic. Very fas ...

12 The Alpine file system



M. R. Brown, K. N. Kolling, E. A. Taft

November 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 4

Publisher: ACM Press

Full text available: [pdf\(2.95 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Alpine is a file system that supports atomic transactions and is designed to operate as a service on a computer network. Alpine's primary purpose is to store files that represent databases. An important secondary goal is to store ordinary files representing documents, program modules, and the like. Unlike other file servers described in the literature, Alpine uses a log-based technique to implement atomic file update. Another unusual aspect of Alpine is that it performs all commu ...

13 A recovery algorithm for a high-performance memory-resident database system



Tobin J. Lehman, Michael J. Carey

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3

Publisher: ACM Press

Full text available: [pdf\(1.50 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

With memory prices dropping and memory sizes increasing accordingly, a number of researchers are addressing the problem of designing high-performance database systems for managing memory-resident data. In this paper we address the recovery problem in the context of such a system. We argue that existing database recovery schemes fall short of meeting the requirements of such a system, and we present a new recovery mechanism which is designed to overcome their shortcomings. The proposed mecha ...

14 Distributed transactions for reliable systems



Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85**, Volume 19 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.44 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

15 File servers for network-based distributed systems



Liba Svobodova

December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(4.23 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

16 Distributed operating systems



Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(5.49 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

17 Multi-level transaction management for complex objects: implementation, performance, parallelism

Gerhard Weikum, Christof Hasse

October 1993 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 2 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(2.83 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Multi-level transactions are a variant of open-nested transactions in which the subtransactions correspond to operations at different levels of a layered system architecture. They allow the exploitation of semantics of high-level operations to increase concurrency. As a consequence, undoing a transaction requires compensation of completed subtransactions. In addition, multi-level recovery methods must take into consideration that high-level operations are not necessarily atomic if multiple pages ...

Keywords: atomicity, complex objects, inter- and intratransaction parallelism, multi-level transactions, performance, persistence, recovery

18 A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems



Walter H. Kohler

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(3.33 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 A comparison of high-availability media recovery techniques



George Copeland, Tom Keller

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data SIGMOD '89**, Volume 18 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.32 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We compare two high-availability techniques for recovery from media failures in database

systems. Both techniques achieve high availability by having two copies of all data and indexes, so that recovery is immediate. "Mirrored declustering" spreads two copies of each relation across two identical sets of disks. "Interleaved declustering" spreads two copies of each relation across one set of disks while keeping both copies of each tuple on separate disks. Both ...

20 XEL: extended ephemeral logging for log storage management



John S. Keen, William J. Dally

November 1994 **Proceedings of the third international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: [pdf\(987.81 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Extended ephemeral logging (XEL) is a more general variation of the ephemeral logging (EL) technique for managing a log of database activity on disk; it does not require a timestamp to be maintained with each object in the database. XEL does not require periodic checkpoints and does not abort lengthy transactions as frequently as traditional firewall logging for the same amount of disk space. Therefore, it is well suited for concurrent databases and applications which have a wide distributi ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

backup transaction log databases checkpoints fault tolerance

Found 13,239 of 171,143

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Practical byzantine fault tolerance and proactive recovery](#)



Miguel Castro, Barbara Liskov

November 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 4

Publisher: ACM Press

Full text available: [pdf\(1.63 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

2 [Recovery Techniques for Database Systems](#)



Joost S. M. Verhofstad

June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Publisher: ACM Press

Full text available: [pdf\(2.32 MB\)](#)
 Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

3 [ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging](#)



C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available: [pdf\(5.23 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to

varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

4 Highly available systems for database applications



Won Kim

March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1

Publisher: ACM Press

Full text available: pdf(2.43 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

5 Principles of transaction-oriented database recovery



Theo Haerder, Andreas Reuter

December 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available: pdf(2.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

6 Fast-Start: quick fault recovery in oracle



Tirthankar Lahiri, Amit Ganesh, Ron Weiss, Ashok Joshi

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(78.85 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Availability requirements for database systems are more stringent than ever before with the widespread use of databases as the foundation for ebusiness. This paper highlights *Fast-Start™ Fault Recovery*, an important availability feature in Oracle, designed to expedite recovery from unplanned outages. Fast-Start allows the administrator to configure a running system to impose predictable bounds on the time required for crash recovery. For instance, fast-start allows fine-gr ...

7 Fast cluster failover using virtual memory-mapped communication



Yuanyuan Zhou, Peter M. Chen, Kai Li

May 1999 **Proceedings of the 13th international conference on Supercomputing**

Publisher: ACM Press

Full text available: pdf(1.45 MB)



Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

8 Impact of timing constraints on real-time database recovery



Jing Huang, Le Gruenwald

- November 1996 **Proceedings of the workshop on on Databases: active and real-time**
Publisher: ACM Press
Full text available:  pdf(465.79 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

9 Understanding fault-tolerant distributed systems



-  Flavin Cristian
February 1991 **Communications of the ACM**, Volume 34 Issue 2
Publisher: ACM Press
Full text available:  pdf(6.17 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

10 An architecture for fault tolerance in database systems



-  Fred J. Maryanski
January 1980 **Proceedings of the ACM 1980 annual conference**
Publisher: ACM Press
Full text available:  pdf(577.74 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An architecture for fault tolerance in a database management system is based upon the concepts of careful replacement and differential files on multiple media with backup copies available. An algorithm for transaction execution that preserves the highest degree of consistency is presented along with an algorithm for the reorganization of the database. The reorganization algorithm merges the differential file with the original database in an on-line fashion. Thus the database is available co ...

11 A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems

-  Walter H. Kohler
June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2
Publisher: ACM Press
Full text available:  pdf(3.33 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 BASE: Using abstraction to improve fault tolerance

-  Miguel Castro, Rodrigo Rodrigues, Barbara Liskov
August 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 3
Publisher: ACM Press
Full text available:  pdf(438.18 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability ...

Keywords: Byzantine fault tolerance, N-version programming, asynchronous systems, proactive recovery, state machine replication

13 Crash recovery for real-time main memory database systems

Jing Huang, Le Gruenwald



February 1996 **Proceedings of the 1996 ACM symposium on Applied Computing**

Publisher: ACM Press

Full text available: pdf(565.80 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: database, real-time, recovery

14 Distributed, object-based programming systems



Roger S. Chin, Samuel T. Chanson

March 1991 **ACM Computing Surveys (CSUR)**, Volume 23 Issue 1

Publisher: ACM Press

Full text available: pdf(2.97 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The development of distributed operating systems and object-based programming languages makes possible an environment in which programs consisting of a set of interacting modules, or objects, may execute concurrently on a collection of loosely coupled processors. An object-based programming language encourages a methodology for designing and creating a program as a set of autonomous components, whereas a distributed operating system permits a collection of workstations or personal computers ...

Keywords: capability scheme, distributed operating systems, error recovery, method invocation, nested transaction, object model, object reliability, object-based programming languages, processor allocation, resource management, synchronization, transaction

15 Performance analysis of recovery techniques



Andreas Reuter

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Publisher: ACM Press

Full text available: pdf(2.47 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Various logging and recovery techniques for centralized transaction-oriented database systems under performance aspects are described and discussed. The classification of functional principles that has been developed in a companion paper is used as a terminological basis. In the main sections, a set of analytic models is introduced and evaluated in order to compare the performance characteristics of nine different recovery techniques with respect to four key parameters and a set of other pa ...

16 Report on the fourth ACM SIGOPS European workshop fault tolerance support in distributed systems



Özalp Babaoğlu

January 1991 **ACM SIGOPS Operating Systems Review**, Volume 25 Issue 1

Publisher: ACM Press

Full text available: pdf(1.76 MB) Additional Information: [full citation](#), [index terms](#)

17 Approaches to fault-tolerant and transactional mobile agent execution---an algorithmic view



Stefan Pleisch, André Schiper

September 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 3

Publisher: ACM Press

Full text available:  pdf(946.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Over the past years, mobile agent technology has attracted considerable attention, and a significant body of literature has been published. To further develop mobile agent technology, reliability mechanisms such as fault tolerance and transaction support are required. This article aims at structuring the field of fault-tolerant and transactional mobile agent execution and thus at guiding the reader to understand the basic strengths and weaknesses of existing approaches. It starts with a discus ...

Keywords: ACID, Byzantine failures, agreement problem, asynchronous system, commit, crash failures, fault tolerance, malicious places, mobile agents, replication, security, transaction

18 [The Recovery Manager of the System R Database Manager](#)



Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.75 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 [BASE: using abstraction to improve fault tolerance](#)



Rodrigo Rodrigues, Miguel Castro, Barbara Liskov

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01**, Volume 35 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.47 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability ...

20 [Performance analysis of checkpointing strategies](#)







Asser N. Tantawi, Manfred Ruschitzka

May 1984 **ACM Transactions on Computer Systems (TOCS)**, Volume 2 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.15 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: database recovery, equicost checkpointing strategy, equidistant checkpointing strategy, error recovery, performance modeling and optimization, rollback recovery, system availability

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

☐ Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((transaction<in>metadata) <and> (log<in>metadata))<and> (checkpoints<in>me

e-mail

Your search matched 8 of 1320520 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)
[New Search](#)

Modify Search

☐ Check to search only within this results set
Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[Select All](#) [Deselect All](#)

- ☐ 1. **Client-based logging for high performance distributed architectures**
Panagos, E.; Biliris, A.; Jagadish, H.V.; Rastogi, R.;
[Data Engineering, 1996. Proceedings of the Twelfth International Conference \(](#)
26 Feb.-1 March 1996 Page(s):344 - 351
Digital Object Identifier 10.1109/ICDE.1996.492182
[AbstractPlus](#) | Full Text: [PDF\(892 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- ☐ 2. **DSDT: Durable Scripts containing Database Transactions**
Salzberg, B.; Tombroff, N.;
[Data Engineering, 1996. Proceedings of the Twelfth International Conference \(](#)
26 Feb.-1 March 1996 Page(s):624 - 633
Digital Object Identifier 10.1109/ICDE.1996.492213
[AbstractPlus](#) | Full Text: [PDF\(1048 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- ☐ 3. **A crash recovery algorithm based on multiple logs that exploits parallelis**
Kumar, A.;
[Parallel and Distributed Processing, 1990. Proceedings of the Second IEEE Sy](#)
9-13 Dec. 1990 Page(s):156 - 159
Digital Object Identifier 10.1109/SPDP.1990.143525
[AbstractPlus](#) | Full Text: [PDF\(348 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- ☐ 4. **Lazy garbage collection of recovery state for fault-tolerant distributed sh**
Sultan, F.; Nguyen, T.D.; Ifode, L.;
[Parallel and Distributed Systems, IEEE Transactions on](#)
Volume 13, Issue 10, Oct. 2002 Page(s):1085 - 1098
Digital Object Identifier 10.1109/TPDS.2002.1041885
[AbstractPlus](#) | [References](#) | Full Text: [PDF\(1399 KB\)](#) IEEE JNL
[Rights and Permissions](#)
- ☐ 5. **Lazy garbage collection of recovery state for fault-tolerant distributed sh**
Sultan, F.; Nguyen, T.D.; Ifode, L.;
[Parallel and Distributed Systems, IEEE Transactions on](#)
Volume 13, Issue 7, July 2002 Page(s):673 - 686
Digital Object Identifier 10.1109/TPDS.2002.1019857

[AbstractPlus](#) | [References](#) | Full Text: [PDF](#)(1287 KB) IEEE JNL
[Rights and Permissions](#)

- ☐ 6. **On failure recoverability of client-server applications in mobile wireless e**
Chen, I.-R.; Baoshan Gu; George, S.E.; Sheng-Tzong Cheng;
[Reliability, IEEE Transactions on](#)
Volume 54, Issue 1, March 2005 Page(s):115 - 122
Digital Object Identifier 10.1109/TR.2004.837518

[AbstractPlus](#) | [References](#) | Full Text: [PDF](#)(216 KB) IEEE JNL
[Rights and Permissions](#)

- ☐ 7. **Message logging: pessimistic, optimistic, causal, and optimal**
Alvisi, L.; Marzullo, K.;
[Software Engineering, IEEE Transactions on](#)
Volume 24, Issue 2, Feb. 1998 Page(s):149 - 159
Digital Object Identifier 10.1109/32.666828

[AbstractPlus](#) | [References](#) | Full Text: [PDF](#)(280 KB) IEEE JNL
[Rights and Permissions](#)

- ☐ 8. **Post-crash log processing for fuzzy checkpointing main memory databas**
Li, X.; Eich, M.H.;
[Data Engineering, 1993. Proceedings. Ninth International Conference on](#)
19-23 April 1993 Page(s):117 - 124
Digital Object Identifier 10.1109/ICDE.1993.344071

[AbstractPlus](#) | Full Text: [PDF](#)(568 KB) IEEE CNF
[Rights and Permissions](#)

Indexed by
 Inspec

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

☐ Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((backup<in>metadata) <and> (database<in>metadata))<and> (checkpoints<in>..."

☐ e-mail

Your search matched 3 of 1320520 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)
[New Search](#)

Modify Search

☐ Check to search only within this results set
Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

 [Select All](#) [Deselect All](#)

- ☐ 1. **Per-user checkpointing for mobility database failure restoration**
Yi-Bing Lin;
[Mobile Computing, IEEE Transactions on](#)
Volume 4, Issue 2, March-April 2005 Page(s):189 - 194
Digital Object Identifier 10.1109/TMC.2005.28
[AbstractPlus](#) | [References](#) | Full Text: [PDF](#)(448 KB) IEEE JNL
[Rights and Permissions](#)
- ☐ 2. **Recovery technique based on fuzzy checkpoint in a client/server databas**
Eui-In Choi; Hae-Chull Lim;
[Computer Software and Applications Conference, 1996. COMPSAC '96., Proc International](#)
21-23 Aug. 1996 Page(s):542 - 547
Digital Object Identifier 10.1109/CMPSAC.1996.545881
[AbstractPlus](#) | Full Text: [PDF](#)(440 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 3. **Restoration scheme of mobility databases by mobility learning and predi networks**
Joon-Min Gil; Chan Yeol Park; Chong-Sun Hwang; Doo-Soon Park; Jin Gon Si Jeong;
[Selected Areas in Communications, IEEE Journal on](#)
Volume 19, Issue 10, Oct. 2001 Page(s):1962 - 1973
Digital Object Identifier 10.1109/49.957311
[AbstractPlus](#) | [References](#) | Full Text: [PDF](#)(208 KB) IEEE JNL
[Rights and Permissions](#)

 Indexed by
[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

Search Results**BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "((backup<in>metadata) <and> (store<in>metadata))<and> (checkpoints<in>meta

e-mail

Your search matched 1 of 1320520 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)[New Search](#)

Modify Search

((backup<in>metadata) <and> (store<in>metadata))<and> (checkpoints<in>meta

Search☐ Check to search only within this results setDisplay Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

view selected items**Select All** **Deselect All**☐ **1. Restoration scheme of mobility databases by mobility learning and prediction networks**

Joon-Min Gil; Chan Yeol Park; Chong-Sun Hwang; Doo-Soon Park; Jin Gon Si Jeong;

[Selected Areas in Communications, IEEE Journal on](#)
Volume 19, Issue 10, Oct. 2001 Page(s):1962 - 1973
Digital Object Identifier 10.1109/49.957311[AbstractPlus](#) | [References](#) | Full Text: [PDF\(208 KB\)](#) IEEE JNL
[Rights and Permissions](#)Indexed by
 Inspec[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

☐ Search Results
[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "((backup<in>metadata) <and> (transaction<in>metadata))<and> (checkpoints<in>metadata))"

Your search matched 2 of 1320520 documents.

e-mail

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)[New Search](#)

Modify Search

☐ Check to search only within this results set
Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[Select All](#) [Deselect All](#)

- ☐ 1. Effect of fault tolerance on response time-analysis of the primary site application
Huang, Y.; Jalote, P.;
[Computers, IEEE Transactions on](#)
Volume 41, Issue 4, April 1992 Page(s):420 - 428
Digital Object Identifier 10.1109/12.135555
[AbstractPlus](#) | Full Text: [PDF\(820 KB\)](#) IEEE JNL
[Rights and Permissions](#)
- ☐ 2. Per-user checkpointing for mobility database failure restoration
Yi-Bing Lin;
[Mobile Computing, IEEE Transactions on](#)
Volume 4, Issue 2, March-April 2005 Page(s):189 - 194
Digital Object Identifier 10.1109/TMC.2005.28
[AbstractPlus](#) | [References](#) | Full Text: [PDF\(448 KB\)](#) IEEE JNL
[Rights and Permissions](#)

 Indexed by
[Help](#) [Contact Us](#) [Privacy & Policy](#)

© Copyright 2006 IEEE –


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

☐ Search Results[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "((backup<in>metadata) <and> (log<in>metadata))<and> (checkpoints<i..."

e-mail

Your search matched 1 of 1320520 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)[New Search](#)

Modify Search

((backup<in>metadata) <and> (log<in>metadata))<and> (checkpoints<in>metad

Search

☐ Check to search only within this results setDisplay Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[view selected items](#)[Select All](#) [Deselect All](#)

- ☐ 1. **Recovery technique based on fuzzy checkpoint in a client/server databas**
 Eui-In Choi; Hae-Chull Lim;
Computer Software and Applications Conference, 1996. COMPSAC '96., Proc
International
 21-23 Aug. 1996 Page(s):542 - 547
 Digital Object Identifier 10.1109/CMPSAC.1996.545881
[AbstractPlus](#) | Full Text: [PDF](#)(440 KB) IEEE CNF
[Rights and Permissions](#)

 Indexed by
 Inspec
[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –



Welcome United States Patent and Trademark Office

☐ [Search Session History](#)
[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Tue, 28 Feb 2006, 11:46:03 AM EST

Edit an existing query or
compose a new query in the
Search Query Display.

Search Query Display



Select a search number (#)
to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- | | |
|----------------------------|--|
| <u>#1</u> | ((backup<in>metadata) <and> (transaction<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#2</u> | ((backup<in>metadata) <and> (transaction<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#3</u> | ((backup<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#4</u> | ((backup<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#5</u> | ((backup<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#6</u> | ((backup<in>metadata) <and> (mergeing<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#7</u> | ((backup<in>metadata) <and> (merging<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#8</u> | ((backup<in>metadata) <and> (sort<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#9</u> | ((backup<in>metadata) <and> (database<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#10</u> | ((backup<in>metadata) <and> (preceding<in>metadata))
<and> (checkpoints<in>metadata) |
| <u>#11</u> | ((backup<in>metadata) <and> (store<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#12</u> | ((transaction<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#13</u> | ((transaction<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#14</u> | ((transaction<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |
| <u>#15</u> | ((transaction<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata) |

#16 ((transaction<in>metadata) <and> (log<in>metadata))<and>
(checkpoints<in>metadata)

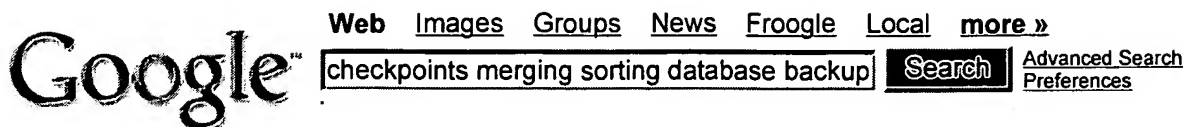


Indexed by



[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –

[Sign in](#)

Web Results 11 - 20 of about 58,900 for checkpoints merging sorting database backup transaction log ide

[PDF] Recovery for Shared Disk Systems Using Multiple Redo Logs

File Format: PDF/Adobe Acrobat - [View as HTML](#)

stored on the **log**. •. TID is a unique **identifier** for the **transaction**. It may help us find the undo ... This safe point is the product of **database backup**, ...
www.hpl.hp.com/techreports/Compaq-DEC/CRL-90-4.pdf - [Similar pages](#)

MySQL Reference Manual for version 3.23.39. - 8 MySQL Table Types

See section 22.2 **Database Backups**. Warning: If you delete old **log** files that ... A 6-byte field tells the **transaction identifier** for the last **transaction** ...
playground.utu.fi/mysql/manual_Table_types.html - 141k - [Cached](#) - [Similar pages](#)

[PDF] Main Memory Database Systems: An Overview

File Format: PDF/Adobe Acrobat

backup copy (see Section I) and to keep a **log** of **transaction**. activity. ... **sorting** (for **sort-merge** joins) is not beneficial when the data ...
doi.ieeecomputersociety.org/10.1109/69.180602 - [Similar pages](#)

[PDF] Backup and Recovery

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The **transaction log** is not truncated when a **database backup** is created. If you plan. to use **database backups** only, the truncate **log** on the **CHECKPOINT** option ...
www.sqlbook.com/codecentric/chapters/ch03.pdf - [Similar pages](#)

Release Notes Release 7.4 Release date: 2003-11-17 ...

It can instead be a non-first character in **identifiers**. This was done to improve compatibility with other **database** systems, and to avoid syntax problems ...
www.dbas.com.ar/postgresql/README.v7.4 - 40k - [Cached](#) - [Similar pages](#)

[PDF] Write-Optimized B-Trees

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Thus, when backing up the **log** to. complement the online **database backup**, migration trans- ... **identifiers** that user **transactions** may temporarily retain. ...
www.vldb.org/conf/2004/RS18P2.PDF - [Similar pages](#)

Englanninkieliset relaatiotietokantatermit

archival **database** (*) archive archive **log** array fetch array select ... **backup** (2) base table basic access batch processing begin **transaction** begin work ...
www.cs.helsinki.fi/relaatiosanasto/rseng_termit.html - 24k - [Cached](#) - [Similar pages](#)

Glossary of Terms

backup. A copy of a **database**, filegroup, file, or **transaction log** that can be ... Character data is **sorted** using collation information, including locale, ...
[msdn2.microsoft.com/en-us/library\(d=robot\)/ms165911.aspx](http://msdn2.microsoft.com/en-us/library(d=robot)/ms165911.aspx) - 209k - [Cached](#) - [Similar pages](#)

Manucorp.com - MySQL Reference Manual for version 4.0.10-gamma ...

All used tables must be in the same **database** as the **MERGE** table itself. ... InnoDB: Starting **log** scan based on **checkpoint** at InnoDB: **log** sequence number 0 ...
www.manucorp.com/docs/mysql/manual_Table_types.html - 101k - Feb 26, 2006 -

[Cached](#) - [Similar pages](#)

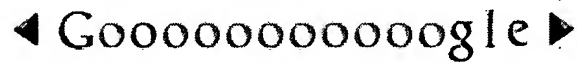
[PDF] The Recovery Manager of the System R Database Manager

File Format: PDF/Adobe Acrobat - [View as HTML](#)

redo are based on records kept in a **transaction log**. The **checkpoint** mechanism is ...

Keywords and Phrases **transactions, database, recovery, reliability** ...

www.scs.cs.nyu.edu/V22.0480-003/sched/recovery.pdf - [Similar pages](#)



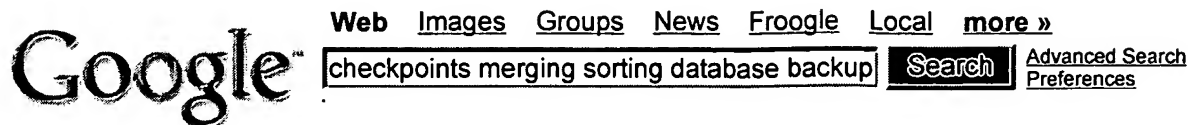
Result Page: **[Previous](#)** [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [11](#) **[Next](#)**

checkpoints merging sorting databas

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)

Web Results 21 - 30 of about 58,900 for checkpoints merging sorting database backup transaction log ide

MySQL Reference Manual for version 5.0.0-alpha. - 14 MySQL Table Types

Before 4.1.1 all used tables had to be in the same **database** as the **MERGE** table ...

InnoDB: Starting **log** scan based on **checkpoint** at InnoDB: **log** sequence ...

www.hupo.org.cn/docs/mysql/manual_Table_types.html - 217k - [Cached](#) - [Similar pages](#)

MySQL Reference Manual for version 4.1.1-alpha. - 7 MySQL Table Types

A 6-byte field tells the **transaction identifier** for the last **transaction** which ... See section

4.5.1 **Database Backups**. Warning: If you delete old **log** files ...

info.ccone.at/INFO/MySQL-Doc/manual_Table_types.html - 204k - [Cached](#) - [Similar pages](#)

Glossary

... stamp of the **database** creation, the current **log** sequence number, and **checkpoint** information. ... A globally unique **identifier** for a row in a **database**. ...

www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96524/glossary.htm - 45k -

[Cached](#) - [Similar pages](#)

[PDF] Main memory database systems: an overview - Knowledge and Data ...

File Format: PDF/Adobe Acrobat

transaction identifier with each **log** record. The processor can ... disk, and **checkpoint** servers keep the. disk-resident **backup database** up ...

dx.doi.org/10.1109/69.180602 - [Similar pages](#)

[PDF] COPYRIGHTED MATERIAL

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Dictionary **sort** order, 25. differential **database backups**, 265,. 275–276, 515 ... to **backup transaction log**, 274–275. to change **database** options, 137–139 ...

media.wiley.com/product_data/excerpt/88/07821427/0782142788-2.pdf - [Similar pages](#)

Englanninkieliset relaatiotietokantatermit

backup (2) base table basic access batch processing begin **transaction** begin work ...

transaction management trigger triggered action tuple tuple **identifier** ...

www.cs.helsinki.fi/relaationsanasto/v1996/rseng_termit.html - 21k - [Cached](#) - [Similar pages](#)

Chapter 3: Error Message Writeups

To accomplish this, the **checkpoint** process periodically checks the number of **log** records in the **transaction log** for each **database**. The number of **log** records ...

infocenter.sybase.com/help/topic/com.

sybase.help.ase_12.5.svrtsq/html/svrtsq/svrtsq195.htm - 61k - [Cached](#) - [Similar pages](#)

[PDF] Implementing a Database in Microsoft SQL Server 7.0 & Visual Basic 6.0

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Server 7.0 **backup** methods. Use the **BACKUP LOG** statement to. manage **transaction**

logs. Design an appropriate **backup** strategy. Module 7: Restoring **Databases** ...

www.f1comp.co.uk/sql/sql7all.pdf - [Similar pages](#)

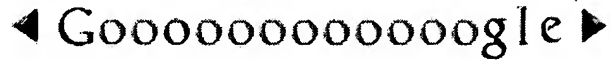
EMBEDDED OBJECT ORIENTED DATABASE SYSTEMS FOR JAVA AND CSHARP

That is why almost all modern DBMS use a **transaction log**. But for the Perst embedded **database**, multi-client access (which requires multiple concurrent ...

www.perst.org/eodbms.html - 111k - [Cached](#) - [Similar pages](#)

SQL Server 2000 versus Oracle

Both data pages and **transaction log** pages are copied to the **backup** set. ... SQL Server uses globally unique **identifier** columns for **merge** replication to ...
www.akadia.com/services/sqlsrv2ora.html - 513k - [Cached](#) - [Similar pages](#)



Result Page: [Previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [11](#) [12](#) [Next](#)

checkpoints merging sorting databas

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

backup database

Search

[Advanced Search](#)
[Preferences](#)**Web**Results 1 - 10 of about **106,000,000** for **backup database** . (0.19 seconds)**BACKUP (Transact-SQL Reference (SQL Server))**

For more information about **database backup** and restore operations, see Backing Up and Restoring ... Syntax Backing up an entire **database: BACKUP DATABASE** { msdn.microsoft.com/library/ en-us/tsqlref/ts_ba-bz_35ww.asp - 47k - [Cached](#) - [Similar pages](#)

BACKUP DATABASE Command

db2 **backup database** sample use tsm open 2 sessions with 4 buffers ... db2 **backup database** payroll tablespace (syscatspace, userspace1) to /dev/rmt0, ... publib.boulder.ibm.com/infocenter/db2help/topic/com.ibm.db2.udb.doc/core/r0001933.htm - 24k - [Cached](#) - [Similar pages](#)

Backup Database Tutorial

Please be patient as file downloads... Orion **Backup Database** Tutorial. www2.solarwinds.net/Tutorials/Backups.html - 2k - [Cached](#) - [Similar pages](#)

Backing up your database « WordPress Codex

It is strongly recommended that you **backup** your **database** at regular ... Restoring your **database** from **backup** is then possible if something goes wrong. ... codex.wordpress.org/Backing_up_your_database - 25k - [Cached](#) - [Similar pages](#)

Backup and restore databases

Article about backing up and restoring **databases** and transaction logs. MSSQLCity.Com - all about MS SQL (SQL Server Articles, FAQ, Scripts, Tips and Test ... www.mssqlcity.com/Articles/Adm/**backup_database**.htm - 21k - [Cached](#) - [Similar pages](#)

Choosing the Right Backup in SQL Server

BACKUP DATABASE database_name FILE = file_name TO DISK = 'file_for_backup' ... **BACKUP DATABASE** database_name FILEGROUP = file_group_name TO DISK ... www.devx.com/getHelpOn/10MinuteSolution/16507/0/page/3 - 83k - [Cached](#) - [Similar pages](#)

Database Backup and Recovery, Database Backup and Recovery

CA's **backup** and recovery solution helps you with day-to-day management and disaster recovery planning. www3.ca.com/Solutions/SubSolution.aspx?ID=2852 - 21k - [Cached](#) - [Similar pages](#)

MySQL 5.0 Reference Manual :: 5.10.1 Database Backups**Sponsored Links****BitLeap Tapeless Backup**

Backup to two secure locations
From \$49/month for 20GB!
www.bitleap.com

Best WAN Accelerator 2006

Reduced **backup** windows.
Maximize high-speed WAN Links
www.Riverbed.com

Data Storage White Papers

Free IT Research Library Focused
On Storage Topics From Top Vendors
www.FindWhitePapers.com

Learn Database Programs

Get a Free Guide
To Online Degrees.
www.DevryUniv.com
Washington, DC (Hagerstown, MD)

Backup to tape or disk

Removable or fixed media
Numerous manufacturers & software
www.sunstarco.com

Database Backup

Get deals on **Database Backup**.
Locate vendors on our directories.
www.business.com

Efficient Mac™ Backups

No drives, disks or complicated
software. No time to waste.
www.backjack.com

Remote Backup Solution?

Online **Backup** starts from \$8.99
Free **database** support, many OS
www.perfectbackup.net

MySQL 5.0 Reference Manual :: 5 **Database** Administration :: 5.10 **Backup** and Recovery ... **Database Backups**. Because MySQL tables are stored as files, ...
www.mysql.com/doc/B/a/Backup.html - 28k - [Cached](#) - [Similar pages](#)

MeetingPlace Web MPWEB SQL Database Backup [Cisco MeetingPlace ...

If you want to export the MPWEB **database** to create a **backup** copy (for ... **BACKUP DATABASE** successfully processed 619 pages in 1.709 seconds (2.962 MB/sec) ...

www.cisco.com/en/US/products/sw/ps5664/

ps5669/products_tech_note09186a008021fdc7.shtml - 21k - [Cached](#) - [Similar pages](#)

Backing up and Restoring a Database

You must have DBA authority to use **BACKUP DATABASE**, whose syntax is as follows: ...

Never make changes to a **backup database**! To read more about validation, ...

www.tldp.org/HOWTO/Sybase-ASA-HOWTO/backup.html - 9k - [Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Go

oooooooooooo

gle

Result Page: 1 2 3 4 5 6 7 8 9 10 [Next](#)

Info when you want it, right on your desktop
Free! [Download Google Desktop](#)

Google

News

New lines of communication
Financial Times 3 hrs ago

Email

Lunch tomorrow?
Mandy MY <.. 11 min ago

61°F Clear - Mount

DJI 10434.87 -84

Type to search


11:22 AM

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied?](#) [Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)

	Web Images Groups News Froogle Local more »
	<input type="text" value="merging backup database checkpoints transac"/> <input type="button" value="Search"/> Advanced Search Preferences

Web Results 1 - 10 of about **105,000** for **merging backup database checkpoints transaction log**. (0.59 seco

SQL Server 2000 Backup and Restore

While in Simple Recovery model, a **database's transaction log** is truncated every time a **CHECKPOINT** is invoked for the **database**. *. **Transaction log backups ...**

www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlbackuprest.mspx - 101k - [Cached](#) - [Similar pages](#)

Exchange Server Soft and Hard Recovery

When the **database** is mounted again, Exchange reads the **checkpoint** file and ... The **backup** program restores **transaction log** files from the **backup** set to this ...

www.microsoft.com/technet/prodtechnol/exchange/2003/insider/exchange-recovery.mspx - 37k - [Cached](#) - [Similar pages](#)

Setting Database Options

The **database backup** should be restored, and any **transaction log backups ...** When a **database** is in **log truncate** mode, an automatic **checkpoint** occurs in the ...

doc.ddart.net/mssql/sql70/8_des_03_8.htm - 25k - [Cached](#) - [Similar pages](#)

Glossary

backup. A **database**, **transaction log**, file, or filegroup in a **database**. A **backup** is made to tape, ... See also **merge** replication, synchronous **transaction**. ...

doc.ddart.net/mssql/sql70/gloss01.htm - 198k - [Cached](#) - [Similar pages](#)

Setting Up a Recovery Storage Group

You can **merge** recovered data into a **database** from a recovery storage group **database ...** Do not set the **transaction log** path to the same path as that of the ...

www.netscum.dk/.../prodtechnol/exchange/guides/UseE2k3RecStorGrps/6166b263-eea6-4389-b8c3-1a47462e4335.mspx - 36k - [Cached](#) - [Similar pages](#)

SQL Server Replication - Problems with massive log files and ...

I'm currently running a complete **database backup** and hopefully that'll clean out the **transaction log** file so we can get a few more months out of it. ...

www.mcse.ms/archive95-2004-6-796070.html - 16k - [Cached](#) - [Similar pages](#)

Glossary

When you make an RMAN **backup** while the **database** is open, however, you do not need to put the tablespaces in **backup** mode. online redo **log ...**

www.oracle-10g.de/oracle_10g_documentation/server.101/b10743/glossary.htm - 66k - [Cached](#) - [Similar pages](#)

Drive Recovery Forum :: View topic - Logs Make Disk Full

Transaction logs generated subsequent to the **backup** can be played into the restored **database**, thus bringing it up to date. If any of the **log** files needed ...

forum.dtidata.com/viewtopic.php?t=88 - 44k - [Cached](#) - [Similar pages](#)

Sams Teach Yourself Microsoft SQL Server 7 in 21 Days > Setting ...

... the **transaction log**, you should **back up** your **database** after performing them. Truncate **Log On Checkpoint**— **Checkpoints** are used to write used data pages ...

www.informit.com/library/content.asp?b=STY_Sql_Server_7&seqNum=38 - 28k - [Cached](#) - [Similar pages](#)

[PDF] Transaction Management Database Recovery Types of failures ...File Format: PDF/Adobe Acrobat - [View as HTML](#)

If **database** has been damaged:. • Need to restore last **backup** copy of **database** and reapply. updates of committed **transactions** using **log** file. ...

www.csc.liv.ac.uk/~valli/Comp302/COMP302-dbrecovery-notes.pdf - [Similar pages](#)**Google Groups results for merging backup database checkpoints transaction log**[merging transaction log files](#) - microsoft.public.exchange.admi ... - Jun 27, 2005[Yet another pair of spanish translation updates](#) - comp.databases.postgresql.patc ... - Sep 14, 2004Try your search again on [Google Book Search](#)

Goooooooooooooogle ►
Result Page: 1 2 3 4 5 6 7 8 9 10 [Next](#)

Info when you want it, right on your desktop
Free! [Download Google Desktop](#)

Google

News

New lines of communication
Financial Times 3 hrs ago

Email

Lunch tomorrow?
Mandy M Y <... 11 min ago

61°F Clear - Mount
DJI 10434.87 -84.°

Type to search


11:22 AM

merging backup database checkpoir [Search](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)


[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

[Advanced Search](#)
[Preferences](#)

Web Results **21 - 30** of about **105,000** for **merging backup database checkpoints transaction log**. (0.19 sec)

Database Journal Forums - MS SQL 7

... **database** options/truncate **log** on **checkpoint** · **CREATING VIEWS** · **Database Backups**
 · Cursor is READ ONLY. ... Need help again · **Transaction log** problem! ...
forums.databasejournal.com/archive/index.php/f-24-p-33.html - 20k -
[Cached](#) - [Similar pages](#)
[\[More results from forums.databasejournal.com \]](#)

Amazon.co.uk: Table of Contents Books: SQL Server 7 Performance ...

How the **Transaction Log** Works Truncating the **Transaction Log Checkpoints**
 Summary ... Tuning OLTP **Transactions** Controlling Data Content **Database Backup** ...
www.amazon.co.uk/exec/obidos/tg/stores/detail/-/books/0735609098/toc - 34k -
[Cached](#) - [Similar pages](#)

Full time MS SQL and OCP Certification courses in London, MS SQL ...

Database Writer (DBWn), **Log** Writer (LGWR), **Checkpoint** (CKPT), System Monitor (SMON), ... Describe the basics of **database backup**, restore and recovery ...
www.london-college.com/full_time_training_courses/ms_sql_ocp_fulltime.aspx - 94k -
[Cached](#) - [Similar pages](#)

Sybase Inc - Previous Certifications

Backup and recovery: production **database backups**, truncating the **log**; recovering from disasters; ... D. Describe **transaction & checkpoints** & write-ahead **log** ...
www.sybase.com/support/education/professionalcertification/adaptiveserverenterprise/previouscertifications - 257k - [Cached](#) - [Similar pages](#)

HP 3000 Manuals

CONSTRAINT **checkpoint** and **log** file space archive **checkpoint** record using to ... **log** files **merging** tables moving **log** files with MOVELOG online **backup** in ...
docs.hp.com/cgi-bin/doc3k/B3621690214.17037/196 - 54k - [Cached](#) - [Similar pages](#)

Features introduced in the various server releases | Oracle FAQ

Fast Start recovery - **Checkpoint** rate auto-adjusted to meet roll forward criteria ... Row-level locking; On-line **database backups**; PL/SQL in the **database** ...
www.orafaq.com/faq/features_introduced_in_the_various_server_releases - 25k -
[Cached](#) - [Similar pages](#)

DBAzone.com: Basic Database Administration Techniques

Execute the ALTER **DATABASE BACKUP** CONTROLFILE TO TRACE; statement on the source ... When Oracle fills a **log** group, a **checkpoint** is fired and Oracle then ...
www.dbazine.com/oracle/or-articles/foot2 - 37k - [Cached](#) - [Similar pages](#)

[PPT] Advanced Topics in Databases

File Format: Microsoft Powerpoint - [View as HTML](#)
 To synchronize with normal **transactions**, the checkpointer writes a record to the **log** after a **checkpoint** is completed. In addition, the **database** is generally ...
ece.ut.ac.ir/.../Data%20Base%20Research%20Group/seminars/Raja-Razavi-Siadati/DB%20Technical%20Report_2.ppt - [Similar pages](#)

[\[PDF\] Microsoft PowerPoint - DI420SN modified.ppt](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Create report of **transaction checkpoints**. Can specify START time. Can specify STOP time ... Vary **database** areas offline or retrieval. **Backup database** areas ... www.ca.com/be/english/past-events/idms-ug/advantage-ca-idms-db-database-recovery.pdf - [Similar pages](#)

[PDF](#) **Database Maintenance Plans**

File Format: PDF/Adobe Acrobat - [View as HTML](#)

transaction log backups as necessary to bring the **database** back to an agreed upon service level. ... include the judicious use of **checkpoints** (that flush ... www.embarcadero.com/resources/tech_papers/DB_Maintenance_6_6.pdf - [Similar pages](#)

◀ Goooooooooooooooooole ▶


Result Page: [Previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [11](#) [12](#) [Next](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)


[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

[Advanced Search](#)
[Preferences](#)

WebResults 1 - 10 of about **152,000** for **merging checkpoints transaction log**. (0.45 seconds)Contents

Recovery using the **log** procedure **checkpoint** facility · Recovery without the **checkpoint** ...
 IMS Message Queue Utilization by **Transaction**, Date report ...
publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.dszos.doc/DRLF9MST02.htm - 38k - [Cached](#) - [Similar pages](#)

EP1387269 Hewlett european software patent - Backup system and ...
merging the transaction log with a preceding **checkpoint** to generate a new **checkpoint**.
 11. A method according to claim 10 comprising the step of sorting the ...
gauss.ffii.org/PatentView/EP1387269 - 42k - [Cached](#) - [Similar pages](#)

Kimberly L. Tripp: Improving *my* SQL skills through your ...
 Q: So does the **transaction log** contain all the information about say an ... So, even if you
 had the equivalent to "truncate **log** on **checkpoint**" this ...
www.sqlskills.com/blogs/kimberly/Trackback.aspx?guid=114e2985-b562-493a-ad0a-7e93273ba132 - 81k - [Cached](#) - [Similar pages](#)

Berkeley DB Reference Guide: Reviewing Berkeley DB **log** files
 Reviewing Berkeley DB **log** files. If you are running with **transactions** and ... **txn_child**,
 Commit a child **transaction**. **txn_ckp**, **Transaction checkpoint**. ...
www.sleepycat.com/docs/ref/debug/printlog.html - 13k - [Cached](#) - [Similar pages](#)

Exchange Server Soft and Hard Recovery
 The function of Restore.env in hard recovery is similar to that of the **checkpoint** file in soft
 recovery. Restore.env defines the range of **transaction log** ...
www.microsoft.com/technet/prodtechnol/exchange/2003/insider/exchange-recovery.mspx -
 37k - [Cached](#) - [Similar pages](#)

Setting Database Options
 When select into/bulkcopy is enabled, the **transaction log** is truncated on each **checkpoint**
 even if trunc. **log** on chkpt is not set to true. ...
doc.ddart.net/mssql/sql70/8_des_03_8.htm - 25k - [Cached](#) - [Similar pages](#)

Sybase: tran **log** dump
 Hi all, i want to know that in following case can I dump **transaction log** I am ... **checkpoint**
 go now noone can run a unlogged **transaction** in DBNAME (nice and ...
www.experts-exchange.com/Databases/Sybase/Q_20761541.html - 56k -
[Cached](#) - [Similar pages](#)

Oracle :: General Info
MERGE uses the SORT_MERGE algorithm to process the NOT IN, ... If the value exceeds
 the actual redo **log** file size, **checkpoints** occur only when switching ...
vsbabu.org/oracle/sect01.html - 29k - [Cached](#) - [Similar pages](#)

PDF Sample Exam -- ANSWER KEY
 File Format: PDF/Adobe Acrobat - [View as HTML](#)
transaction processing. Assume that simple **checkpoints** are used. The crash recovery
 algorithm reads the most recent **checkpoint** information from the **log**, ...
www.cise.ufl.edu/~jhammer/classes/4720-5725/SP03/sample_final1.pdf - [Similar pages](#)

www.mcse.ms/archive95-2004-6-796070.html - 16k - [Cached](#) - [Similar pages](#)

Goooooooooooooogle ►

Free! Get the Google Toolbar. Download Now - About Toolbar

Google   Search  377 blocked  Check  AutoLink  AutoFill

merging checkpoints transaction log

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

We are experiencing some difficulties with our database and, consequently, some document links may be incorrect. Please bear with us while we fix the problem.

25 citations found. Retrieving documents...

K. Salem and H. Garcia-Molina. *System M: A transaction processing testbed for memory resident data*. IEEE Transactions on Data and Knowledge Engineering, 2(1):161-172, Mar 1990.

CiteSeer [Home/Search](#) [Document Not in Database](#) [Summary](#) [Related Articles](#) [Check](#)

This paper is cited in the following contexts:

Improving Predictability of Transaction Execution Times.. - Rajeev Rastogi Seshadri (2000) (3 citations) (Correct)

....fast and measured in a few hundreds of nanoseconds. **Thus, storing the entire database in main memory can result in short and bounded transaction execution times. While disk based databases exhibit improved performance if the entire database can fit in the main memory buffer cache, a MMDB (e.g. SGM90, LSC92, JLR 94, DKO 84) improves performance further by dispensing with the buffer manager, and tuning algorithms to the flat storage hierarchy and the reduced cost of indirection.** Also, MMDB schemes attempt to minimize space usage, of vital importance since main memory remains about one

K. Salem and H. Garcia-Molina. *System M: A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161--172, March 1990.

Fault Resilience In Main-Memory Databases: Handling Process.. - Bohannon (1999) (Correct)

....of storage managers for disk resident data. **These** include the storage managers of Exodus [15] and Starburst [36] With the exception of the Starburst main memory storage component [36] we are not aware of any storage manager that is tailored for mainmemory resident data. **System M [64] is a transaction processing test bed for memory resident data, but is not a full feature storage manager.** 34 The Dal i system [41] implemented at Bell Laboratories, is a storage manager for persistent data whose architecture has been optimized for environments in which the database is

....Checkpointing ensures that only a final portion of the log is needed for recovery. To implement checkpointing in Dal i, two copies of the database image are stored on disk, and alternate checkpoints write dirty pages to alternate copies. **This strategy is called ping pong checkpointing (see, e.g. [64]) The ping pong checkpointing strategy permits a checkpoint that is being created to be temporarily inconsistent; i.e. updates may have been written out without corresponding undo records having been written.** However, after writing out dirty pages, sufficient redo and undo log information is

K. Salem and H. Garcia-Molina. *System M: A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161--172, March 1990.

An Almost-Serial Protocol for Transaction Execution in.. - Blott, Korth (2002) (1 citation) (Correct)

....locks had a significant impact on performance. **This** argues for the use of relatively coarser grain locks. **Indeed, the observation that fine grain locking can incur excessive overhead in certain cases was made in the early days of research in lock protocols [12] Several authors, including [8, 9, 10, 18, 20], have noted this in the context of main memory database systems.** Salem and Garcia Molina have proposed an approach where most (short) transactions are scheduled serially, but long transactions and checkpointing operations execute concurrently [20] Priority is generally given to short,

....[12] Several authors, including [8, 9, 10, 18, 20] have noted this in the context of main memory database systems. **Salem and Garcia Molina have proposed an approach where most (short) transactions are scheduled serially, but long transactions and checkpointing operations execute concurrently [20].** Priority is generally given to short, real time transactions. **However,** unlike the algorithm proposed here, Salem and Garcia Molina s algorithm requires locks to be obtained by all classes of transaction, and is also subject to deadlock. **Deadlock,** though typically a non issue from a real world

[Article contains additional citation context not shown here]

K. Salem and H. Garca-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Trans. on Knowledge and Data Engineering, 2(1):161-172, Mar. 1990.

Incremental Garbage Collection of an Almost Unidirectional.. - Oksanen (2000) (Correct)

....causes more writing, and the relative importance of asynchrony grows. **We** have also implemented asynchronous reading. It almost halves recovery times when using two disks, but additional disks bring little further improvement. **In conjunction with their transaction processing testbed, System M [21], Garcia Molina and Salem informally de ned a main memory database benchmark.** After being generously allowed to peek at their source code we were able to reproduce the benchmark. **The** benchmark simulates a credit card application with four tables of varying sizes and eight transaction types, most

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Data and Knowledge Engineering, 2(1):161-172, Mar 1990.

Database Architectures - Delis, Kanitkar, Kollios (1998) (Correct)

....the second bit. **The** benefits of such a scheme rest with the fact that, often in MMDs, records are locked for a short period of time and are released soon after the update. If there is no need to access the hash table frequently, this technique presents an acceptable locking alternative. **System M [70] features an exclusive shared locking scheme with conversion capability from shared to exclusive mode at the segment level (set of records)** 2.1.4 Logging and Commit Protocols Logging is mandatory as the MMD should be able to avoid lost data and or transactions due to media failure. **As** logging

....media failure. **As** logging is the only operation that has to deal with an external device in MDDs, it can become a bottleneck that may adversely affect system throughput. **A number of solutions have been suggested to solve this problem and are based around the concept of a stable main memory space [26, 70, 56, 54, 45].** Whenever a transaction is ready to commit, the transaction writes its changes into stable memory (non volatile RAM) Stable memory is often used to carry the transaction log and can greatly assist in decoupling persistence from atomicity. **Writing** to such a stable log is a fast 11 operation as

[Article contains additional citation context not shown here]

K. Salem and H. Garcia-Molina. System M: *A Transaction Processing Testbed for Memory Resident Data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161--172, March 1990.

Programming over a Persistent Data Space - I (1992) (Correct)

....persistent data. **One** frequently encounters definitions of persistence in terms of a persistent recording medium, or in terms of its long life. **While we intuitively think of persistent data as residing in files on disk or tape storage, rather than RAM, the introduction of memory resident databases [SaG90, Son89] clearly forces a re evaluation of such an intuitive definition.** Persistent data is often long lived, but it need not be. **And** long lived data need not fit our sense of persistent data; the local variables of a nonterminating process such as an airline reservation system can have a very long life.

K. Salem and H. Garcia-Molina, System M: *A Transaction Processing Testbed for Memory Resident Data*, IEEE Trans. on Knowledge and Data Engineering 2,1 (Mar. 1990), 161-172.

The Role of Distributed Shared Memory in Future Experimental.. - Fleisch (Correct)

....a more attractive paradigm as 64 bit address space architectures emerge and the potential of distribution of persistent data storage throughout a cluster of networked machines can be realized by using DSM. **Our group at the University of California is exploring a modified version of System M[8] that operates in our DSM cluster.** Reliability is a serious concern for DSM systems since we believe that much future research work will be directed towards MMDB s that must store persistent data. **For** future DSM systems to be successful, the issue of reliability must be addressed in substantive

K. Salem and H. Garcia-Molina, System M: *A Transaction Processing Testbed for Memory Resident Data*, IEEE Transactions on Knowledge and Data Engineering, Vol. 2, No. 1, March 1990, pp. 161-172.

Distributed Multi-Level Recovery in Main-Memory Databases - Rastogi, Bohannon.. (1998) (3 citations) (Correct)

....to disk. **As a result of not obtaining latches on pages during updates, it is not possible to enforce the write ahead logging policy, since pages may be updated even as they are being written out. Instead, our recovery algorithm makes use of a strategy called ping pong checkpointing (see, e.g. [19]) In ping pong checkpointing two copies of the database image are stored on disk, and alternate checkpoints write dirty pages to alternate copies.** Writing alternate checkpoints to alternate copies permits a checkpoint that is being created to be temporarily inconsistent; i.e. updates may have

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161-172, 1990.

Distributed Multi-Level Recovery in Main-Memory Databases - Rastogi, Bohannon.. (1998) (3 citations) (Correct)

....for checkpointing, per transaction logs in memory to reduce contention on the system log tail, and recovery using only a single pass over the system log. **The recovery scheme used in Dal i provides several further extensions, such as multi level recovery (20, 14, 13] and fuzzy checkpointing [18, 8].** The goal of the work described here was to extend the Dal i recovery scheme to the distributed memory case, simultaneously maintaining the advantages of the single site scheme, and efficiently supporting the applications described above. **For example, we can make use of transient undo logging to**

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161-172, March 1990.

Dal'i: A High Performance Main Memory Storage Manager - Jagadish Daniel (1994) (7 citations) (Correct)

....transaction processing applications but with much lower latency and higher throughput requirements. **In a typical Dal i application, transactions are small, multiple processes may access shared data and high concurrency, especially on index structures, is important. As a result, 2 System M [SGM90] is a transaction processing test bed for memory resident data, but is not a full feature storage manager.** Dal i supports item level locking. **Also, the recovery algorithm used in Dal i is designed to work well with small transactions.** 3 The Dal i System Architecture A Dal i storage management

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Trans. Knowledge and Data Engineering, 2(1):161-172, March 1990.

Multi-Level Recovery in the Dal'i Main-Memory .. - Bohannon, Rastogi, ... (Correct)

....use. **Thus, data access using a main memory database is very fast compared to using disk based storage managers, even when the disk based manager has sufficient memory to cache all data pages. A number of recovery algorithms for main memory environments have been proposed in the literature [Hag86, SGM90, JSS93, DKO 84, LE93, Eic89, GE91, LC87a, LS92b]** The proposed schemes attempt to reduce logging and flushes to disk by exploiting the fact that in a main memory database, pages are not flushed to disk during normal processing (as in disk based systems) and logs for active transactions fit

....GE91, LC87a, LS92b] The proposed schemes attempt to reduce logging and flushes to disk by exploiting the fact that in a main memory database, pages are not flushed to disk during normal processing (as in disk based systems) and logs for active transactions fit in main memory. **For example, in [SGM90, JSS93] recovery schemes use ping pong checkpointing in which two database images are maintained on disk and successive checkpoints write to alternate copies.** This eliminates the need for writeahead logging (WAL) and reduces the number of log flushes. **Furthermore, in [JSS93] each active**

[Article contains additional citation context not shown here]

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, March 1990.

PPOST: A Parallel Database in Main Memory - Böszörményi, Eder, Weich (1994) (Correct)

....log machine. The log machine would ideally store the log tail in stable main memory. In this case, transactions whose log information arrived in the log machine can be committed immediately. We do not insist, however, on the existence of a stable main memory. **In the lack of this, we precommit [5, 7] the corresponding transactions and let run other transactions (locks are released) In the meantime, the log information is stored on disc in the form of simple sequential files (this can be done at full disc speed) After that, precommitted transactions may be committed.** In case of a system

H. Garcia-Molina, K. Salem, System M: *A Transaction Processing Testbed for Memory Resident Databases*, IEEE Transactions On Knowledge And Data Engineering, Vol. 1, No. 2, March 1990.

Logical and Physical Versioning in Main Memory Databases - Rajeev Rastogi (1997) (7 citations) (Correct)

....of the Twenty Third International Conference on Very Large Data Bases, August 26 29, Athens, Greece, Morgan Kauffman. **1 Introduction While disk based databases exhibit improved performance if the entire database can fit in the main memory buffer cache, a main memory database (MMDB) e.g. SGM90, JLR 94] improves performance further by dispensing with the buffer manager, and tuning algorithms to the flat storage hierarchy and the reduced cost of indirection.** Also, MMDB schemes attempt to minimize space usage, of vital importance since main memory remains about one hundred times as

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, March 1990.

Moving Distributed Shared Memory to the Personal Computer.. - Fleisch, Hyde, Juul (1993) (1 citation) (Correct)

....database, and a battleship simulation [Novak 91, Rawdon 91] Furthermore, most of the applications have been instrumented so that they may perform timing and logging of memory access patterns. 3. **1 Memory Resident Database We are currently implementing a memory resident database based on System M [Salem 90] Memory resident databases hold much promise when used in conjunction with Distributed Shared Memory.** The original version of System M was operational under the Mach operating system and needed substantial modifications to operate in our environment. **Activities** to port system M were broken up

Kenneth Salem and Hector Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161– 172, March 1990.

The Role of Self-Regulation in Corporate Governance: .. - de Jong, DeJong.. (2000) (Correct)

....impact of memory resident databases on concurrency control, commit processing, access methods, data representation, query processing, recovery, performance issues, the application programming interface, and protection. **Prototypes of several memory resident database system have been described in [5, 24, 23, 12, 14, 15, 13].** 10] introduces the use of client caches to reduce transaction response time in a disk based database system. **Several** papers have investigated methods for using underutilized CPU capacities of workstations [18, 20, 8, 6, 7] or load balancing for distributed database applications [25] Many

K. Salem and H. Garsia-Molina. System M: *A Transaction Processing Testbed for Memory Resident Data*. IEEE Transaction on Knowledge and Data Engineering, March 1990.

Transactions for Amadeus - Taylor (1993) (1 citation) (Correct)

....facility for user defined abstract types. **Other** transaction systems were developed on emerging micro kernel based operating systems. **These** systems took advantage of the virtual memory facilities provided by the micro kernels. **Example** of such systems are Camelot (see section 2. 3) and **System M [21] both implemented on the Mach [22] operating system.** System M was designed as a transaction facility for memory resident databases.

Another limitation identified in [18] was that transactions were not unified with programming languages. **Many** projects have addressed this problem, most notably Argus

....independent of all languages. **One** way to perform access detection in the kernel is to use virtual memory techniques. **Some modern operating systems (e.g. Mach [22] and Chorus [57] provide advanced virtual memory facilities which have been used as the basis for access detection and recovery [21, 58].** Using these facilities obviously places a dependence on the underlying operating system. **The** Unix system (on which Amadeus is implemented) provides facilities for protecting an area of virtual memory against reads and writes. **There** are three main disadvantages of using virtual memory techniques

Kenneth Salem and Hector Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, March 1990.

The Architecture of the Dalí Main-Memory Storage Manager - Bohannon, al. (1997) (Correct)

....since pages may be updated even as they are being written out. **For** correctness, in the absence of write ahead logging, two copies of the database image are stored on disk, and alternate checkpoints write dirty pages to alternate copies. **This strategy is called ping pong checkpointing (see, e.g. [40]) The ping pong checkpointing strategy permits a checkpoint that is being created to be temporarily inconsistent; i.e. updates may have been written out without corresponding undo records having been written.** However, after writing out dirty pages, sufficient redo and undo log information is

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, 1990.

The Architecture of the Dalí Main-Memory Storage Manager - Bohannon, al. (1997) (Correct)

....for reducing such costs. **Thus**, they provide an excellent motivation for closely examining the system design of a mainmemory database and tuning it to remove bottlenecks, and have thereby influenced our work significantly. **Much of the work on main memory databases has concentrated on recovery [13, 19, 30, 31, 39].** The work by Eich [13] provides a survey, and the performance studies using System M by Salem and Garcia Molina [39] provide both a good review and performance comparison of many of the schemes suggested by earlier work. **Our** recovery algorithm is in many ways similar to the fuzzy checkpointing

....and tuning it to remove bottlenecks, and have thereby influenced our work significantly. **Much of the work on main memory databases has concentrated on recovery [13, 19, 30, 31, 39] The work by Eich [13] provides a survey, and the performance studies using System M by Salem and Garcia Molina [39] provide both a good review and performance comparison of many of the schemes suggested by earlier work.** Our recovery algorithm is in many ways similar to the fuzzy checkpointing schemes of [39] including use of ping pong checkpointing and dirty page bits. **One** difference is that updates in

[Article contains additional citation context not shown here]

K. Salem and H. Garcia-Molina. System M: *A transaction processing testbed for memory resident data*. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, March 1990.

Incremental Recovery In Main Memory Database Systems - Levy, Silberschatz (1992) (25 citations) (Correct)

....updates to backup database on disk, since paging is expected to be a relatively rare activity. **Many recent research efforts go to the extreme with this trend arguing that there are cases where the entire database can fit in memory, thus eliminating paging entirely (e.g. DKO 84, LN88, GMS90, SGM87a]** With infrequent page replacements, checkpointing and keeping a stable copy of the database may become a very disruptive function. Typically, persistence and atomicity of transactions is guaranteed by performing disk I/O at certain critical points (e.g. flushing a commit log

....dumping to the disk is the basis to all of them. **Therefore**, any variation of direct checkpointing is bound to delay transaction processing to a considerable extent. **Many of the proposed algorithms and schemes for MMDBs rely on the explicit assumption that the entire database is memory resident [GMS90, LN88, SGM87b, DKO 84]** Although other proposals acknowledge that this assumption is not valid for practical reasons, the issue is not addressed directly in their designs [LC87, Hag86, Eic86] Even though the size of

main memory is increasing very rapidly the size of future databases is

H. Garcia-Molina and K. Salem. System M: A transaction processing testbed for memory resident data. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, 1990.

Recovering from Main-Memory Lapses - Jagadish, Silberschatz, Sudarshan (1993) (14 citations) (Correct)

....to the checkpoint location pointer on stable store. **After** this, the old checkpoint may be deleted. **We** assume that there is a pointer in stable store to the latest checkpoint. **The** last action performed during a checkpoint is the update of this pointer. **Thus, we follow a ping pong scheme (see [17]) keeping up to two copies of the database.** Partially written checkpoints are ignored in the event of a crash, and the previous (complete) checkpoint is used instead, so the writing of the checkpoint is atomic (i.e. it either happens completely or appears to have not happened at all) The set of

....Moreover, replaying could become a bottleneck, since it is in effect replaying the committed actions of the mainmemory database on the disk database, in serialization order, and could require a considerable amount of I O. **The black white checkpointing technique of Pu [15] also described in [17], allows action consistent checkpointing, but either requires deferred updates (shadow paging) or undo logging on disk.** The disadvantages of requiring deferred updates were discussed in Section 1. **Hagmann [10] allows fuzzy checkpointing, that is, does not even require action consistency. However,**

K. Salem and H. Garcia-Molina. System M: A transaction processing testbed for memory resident data. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–172, 1990.

An Overview of Real-Time Database Systems - Kao, Garcia-Molina (1995) (25 citations) Self-citation (Garcia-molina) (Correct)

No context found.

H. Garcia-Molina, K. Salem: System M: A Transaction Processing Testbed for Memory Resident Data. IEEE Transactions on Knowledge and Data Engineering. #March 1990# Vol. 2, number 1 161#172

Strip: A Soft Real-Time Main Memory Database for Open Systems - Adelberg (1997) (1 citation) Self-citation (Garcia-molina) (Correct)

....and to speed processing. **The DBGraph model described by [PTV90] makes even more use of pointers and suggests some interesting future extensions to STRIP. Memory** residency of data has implications on all of the major components of a DBMS. **The impact on recovery has been been thoroughly studied and [SGM92] provides an excellent overview.** Main memory access methods and join algorithms must also be modified. **For** instance tree indices can be much deeper since pointer traversal is cheap [LC86b] and join algorithms based on sequential scans (e.g. sort merge) are less attractive [LC86a] An interesting

K. Salem and H. Garcia-Molina. System m: A transaction processing testbed for memory resident data. IEEE Transactions on Knowledge and Data Engineering, 2(1):161–72, 1992.

An Overview of Real-Time Database Systems - Kao, Garcia-Molina (1994) (25 citations) Self-citation (Garcia-molina) (Correct)

....and often unpredictable disk access delays. **As** the price of memory continues to drop, one possible remedy is to put data directly into memory, thus eliminating I O accesses. **In** this subsection, we give a brief account on memory resident database system design. **Interested readers are referred to [10, 18, 19, 47] for further reading.** Compared to disk, main memory access time is much faster (1000 to 10000 times) and is more predictable (no disk seek) These features are very desirable in RTDBSs, and may even be necessary if transactions have extremely tight time constraints. **However, putting all the data**

H. Garcia-Molina, K. Salem: System M: A Transaction Processing Testbed for Memory Resident Data. IEEE Transactions on Knowledge and Data Engineering (Mar. 1990), Vol. 2, number 1 161–172.

An Overview of Real-Time Database Systems - Kao, Garcia-Molina (1995) (25 citations) Self-citation (Garcia-

molina) [\(Correct\)](#)

No context found.

1--8 19. H. Garcia-Molina, K. Salem: *System M: A Transaction Processing Testbed for Memory Resident Data*. IEEE Transactions on Knowledge and Data Engineering.

[A Replicated and Persistent Functional Programming Environment -.. - Oksanen \(2001\)](#) [\(Correct\)](#)

No context found.

K. Salem and H. Garcia-Molina. *System M: A transaction processing testbed for memory resident data*. IEEE Transactions on Data and Knowledge Engineering, 2(1):161-172, Mar 1990.

[Online articles have much greater impact](#) [More about CiteSeer.IST](#) [Add search form to your site](#) [Submit documents](#) [Feedback](#)

CiteSeer.IST - Copyright [Penn State](#) and [NEC](#)